
A Passage-of-time Model of the Cerebellar Purkinje Cell

Matthew Ricci
Brown University
matthew_ricci_1@brown.edu

Junkyung Kim
Brown University
junkyung_kim@brown.edu

Fredrik Johansson
Lund University
fredrik.johansson@med.lu.se

Abstract

The cerebellar Purkinje cell controlling eyeblinks can learn, remember and reproduce the interstimulus interval in a classical conditioning paradigm. Given temporally separated inputs, the cerebellar Purkinje cell learns to pause its tonic inhibition of a motor pathway with high temporal precision so that an overt blink occurs at the right time. Most models relegate the Purkinje cell's passage-of-time representation to afferent granule cells, a subpopulation of which is supposedly selected for synaptic depression in order to make the Purkinje cell pause. However, granule cell models have recently faced two crucial challenges: 1) bypassing granule cells and directly stimulating the Purkinje cell's pre-synaptic fibers during training still produces a well-timed pause, and 2) the Purkinje cell can reproduce the learned pause, invariant to the temporal structure of probe stimulation. Here, we present a passage-of-time model which is internal to the Purkinje cell and is invariant to probe structure. The model accurately simulates the Purkinje cell learning mechanism and makes testable electrophysiological predictions. Importantly, the model is a numerical proof-of-principle for a biological learning mechanism which does not rely on changes of synaptic strength.

1 Introduction

While long-term potentiation (LTP) and long-term depression (LTD) are often measurable during memory encoding, the prevailing doctrine that such Hebbian changes in synaptic strength entirely explain memory has been challenged by several recent findings. In both *Aplysia* (Chen et al., 2014) and mammalian hippocampal cells (Ryan et al., 2015), long-term memory does not in all cases require the persistence of changes in synaptic strength. The result that deviates the furthest from Hebbian theory comes from a recent result in eyeblink conditioning, which showed that a single Purkinje cell can memorize a temporal relationship of hundreds of milliseconds between two inputs (Johansson et al., 2014).

If a conditional stimulus (CS), such as a tone, is repeatedly followed by a blink-eliciting unconditional stimulus (US), with a fixed temporal delay (the interstimulus interval or ISI), a blink response to the CS develops. This conditioned response (CR) occurs just before the expected US onset for ISIs from ~ 100 ms to seconds (Gormezano and Moore, 1969; Kehoe and Macrae, 2002). The underlying learning occurs in a specific microzone within the C3 zone of cerebellar cortical lobule HVI (Yeo et al., 1985; Hesslow, 1994b,a; Mostofi et al., 2010; Heiney et al., 2014). Because many and diverse conditional stimuli may predict the

same unconditional stimulus, the association-forming mechanism should exhibit extensive fan-in, which the Purkinje cell indeed does; its vast dendritic arbor gets CS input from approximately 200,000 parallel fibers (Harvey and Napper, 1991) that come from the tiny and densely packed granule cells (GRC) in the granular layer of the cerebellar cortex, while it gets US input from a single climbing fiber (Mauk et al., 1986; Steinmetz et al., 1989; Hesslow et al., 1999; Hesslow and Yeo, 2002) (Fig. 1) that contacts the entire arbor (Eccles et al., 1967).

Purkinje cells, which are the sole output of the cerebellum, have high tonic firing rates due to an




Figure 1: *Simplified neural circuitry.* The CS is transmitted via the mossy/parallel fiber system and the US via the climbing fiber. Key: IO: inferior olive. cf: climbing fiber. mf: mossy fibers. Pc: Purkinje cell. Gc: Golgi cell. pf: parallel fibers. GRC: granule cells. DCN: deep cerebellar nuclei.

internal pacemaker mechanism (Cerminara and Rawson, 2004). Given CS-US pairings, blink-controlling Purkinje cells learn an adaptively timed pause response (Jirenhed et al., 2007; Halverson et al., 2015; Hesslow, 1995; Hesslow and Ivarsson, 1994; Jirenhed and Hesslow, 2011a) in their tonic inhibition of a motor pathway (Hesslow and Yeo, 2002). This learned pause mirrors the known features of behavioral CRs. It is extinguished during repeated CS alone presentations and it is rapidly re-acquired (Jirenhed et al., 2007). The pause is always adaptively timed, reaches its maximal amplitude just before the predicted onset of the US, and ends shortly after, even if the CS lasts only a few milliseconds or outlasts the ISI by hundreds of milliseconds (Jirenhed and Hesslow, 2011b; Johansson et al., 2014). In temporal uncertainty paradigms where mixed trials of different ISIs are used, multiple temporally locked pause responses to the same CS are learned (Halverson et al., 2015). Importantly, the behavioral lower limit of a minimal ISI of ~ 100 ms holds at the level of the single Purkinje cell (Wetmore et al., 2014).

Contemporary modeling is dominated by the idea of the granular layer network generating a passage-of-time (POT) representation. (Buonomano and Mauk, 1994; Medina and Mauk, 2000; Yamazaki and Tanaka, 2009). Before learning, the Purkinje cell is assumed to fire at its tonic rate due to net excitation from the balanced activity of excitatory granule cells and inhibitory basket/stellate interneurons (Medina et al., 2000).

In response to a CS drive from the mossy fibers, time-variance in granule cell (GRC) activity is assumed to arise from fast and partly lateral feedback inhibition from Golgi cell interneurons. The effect would be a series of random transitions between activity and quiescence in each granule cell, creating instantaneous population vectors that represent time. Spike-timing dependent plasticity then selectively recruits LTD for those GRC-to-Purkinje cell synapses most activated by the CS around the time of US onset. The parallel fiber-Purkinje cell synapse is anti-Hebbian in that correlated excitatory inputs from parallel fibers and the climbing fiber causes LTD. LTD tips the balance of afferent activity towards net inhibition near expected US onset, so that, after learning, the CS alone initiates a well-timed pause in Purkinje cell firing.

Granule cell POT models cannot readily be reconciled with two recent results, one physiological and one computational. On the physiological side, Johansson and colleagues showed that the Purkinje cell itself learns the ISI (Johansson et al., 2014). They bypassed the granular layer network by delivering the CS directly to the parallel fibers. Purkinje cells learned well-timed pause responses to ISIs of 150, 200 and 300 ms. Furthermore, both GABA-ergic interneurons Johansson et al. (2014) and glutamatergic AMPA receptors Johansson et al. (2015) could be blocked without disrupting the pause response. Their result implies that the precisely timed Purkinje cell activity depends upon an internal timing mechanism that measures and stores temporal duration.

The computational process embodied by this internal mechanism deviates significantly from that of previous POT models. After training, the Purkinje cell can produce a well-timed pause given only the initial (< 20 ms) part of the CS (Svensson and Ivarsson, 1999; Jirenhed and Hesslow, 2011b). For example, Johansson et al. (2014) found that drastically varying the CS on probe trials (17.5-800 ms, 100-400 Hz) had no effect on the learned pause. Existing POT models, on the other hand, require the probe stimulus to endure for at least as long as the training stimulus. For, as soon as the probe terminates, afferent activity would shift to its pre-stimulation conditions and return the net input of the Purkinje cell to excitation, preventing pause expression. The real computational process underlying the CR must contain a strong non-linearity which switches the Purkinje cell from a tonic spiking state to a pause-expressing state.

Hence, we see two crucial criteria for the modeling of a Purkinje cell timing mechanism:

Physiological: The neural machinery for pause learning and expression must be contained in the Purkinje cell itself.

Computational: The computational process initiating pause expression must be gated by a strong non-linearity or “switch”.

This second criterion has new neurobiological underpinning in the work of Johansson et al. (2015), who found that nanoliter injections of antagonists of the inhibitory metabotropic glutamate receptor 7 (mGluR7), targeted at the Purkinje cell dendrites, eliminates the pause response. This suggests that mGluR7 elicits the response through activation of protein-activated, hyperpolarizing K^+ channels in the family Kir3 or through inactivation of cAMP or Ca^{2+} signaling. In other words, pause expression may be gated by a metabotropic switch.

Though this hypothesis may explain what triggers pause expression, it does not explain the temporal structure of the pause itself or how this structure is learned. An actual trained Purkinje cell receiving a parallel fiber signal, produces a pause in firing with a slightly delayed onset, accurately timed maximum and critically timed offset. Johansson and Hesslow (2014) provided a conceptual theory for how such a timing mechanism could work. They suggested a selection mechanism whereby CS onset prompts the Purkinje cell to release a batch of evolving “recorder units”. These units could be second messengers, some unknown molecular switches, polynucleotides or evolving proteins. The precise biophysical identity of recorder units does not substantially affect the theory of Johansson and Hesslow (2014); the only crucial feature of these units is that they evolve through different time-encoding states beginning at CS onset. Each unit contributes to inhibition of the cell at the interval that it encodes, and the strength of inhibition at any given time since CS is encoded by the quantity of units that encode the same interval, which increases via repeated exposure to CS-US pairs. Recorder units thus constitute an intracellular POT representation. Importantly, this model satisfies the two above criteria: the pause mechanism is internal to the Purkinje cell and it is mediated by a glutamatergic switch.

Below, this theory is instantiated as a computational model with two parts: a “write” module, which

intracellularly records the interstimulus interval and a “read” module, which translates the archived intervals into a timed hyperpolarization event. These states become active only when afferent activity switches the cell from its passive OFF state to its active ON state. We trained this model cell with several paradigms typical to the eyeblink conditioning literature and found:

- Given ISIs ranging from 100 ms to 500 ms, the cell learns a pause response with critically timed onset, maximum and offset, as in Johansson et al. (2014).
- After training, the cell produces the learned pause given only the initial part of the CS, as in Svensson and Ivarsson (1999), Jirenhed and Hesslow (2011b) and Johansson et al. (2014).
- The learned pause is extinguished by CS-only trials in about as much time as initial acquisition, as in Jirenhed et al. (2007); Johansson et al. (2014).
- Alternating ISIs across trials teaches the cell two pauses, as in Jirenhed et al. (2007); Johansson et al. (2014).

Additionally, we ran two simulations, which, to our knowledge, have not been tested on the biological Purkinje cell. First, we tried stimulating the cell with two-part CSs or two-part USs on each trial: stimuli consisted of two impulse trains to the same fiber separated by a period of silence, so that the cell was effectively exposed to two ISIs. In these simulations, our model cell only learns one ISI. Next, we tried varying the intertrial interval (ITI) along with the ISI. We found that time to pause acquisition in all simulations was a function of both of these temporal quantities. The behavior of this formal model may help to inform future electrophysiology.

2 Model

Model structure. The Purkinje cell has two computational goals. First, it must record or “write” an ISI into memory. Second, it must “read” this ISI into a well-timed pause. As stated above, the Purkinje cell only transitions from its passive spiking behavior to its active read/write behavior when some type of switch is activated. Thus, we may formally model the cell as having two modules (read/write; Fig. 2, top vs. bottom row) which can be in two states (OFF/ON; Fig. 2, left vs. right column).

Before CS stimulation, both the read and write modules are OFF and the cell spikes at its tonic rate. When pre-synaptic spikes arrive at the cell, a quantity called “activation energy” (AE_{write} and AE_{read}) begins to increase. When AE_{write} reaches a threshold value, the write module switches to the ON state. A refractory timer on this switch is then reset so that the cell is insensitive to the CS for a given period.

When the write module is activated after sufficient CS stimulation, a batch of “recorder units” is released from a “reserve”. These units are a formal abstraction meant to stand for whatever process encodes the POT within the Purkinje cell. The reserve is a way of imagining these units before the write module is activated. After some recorder units have been released from the archive, the reserve begins to replenish at a constant rate. The released recorder units evolve through states in a way that mirrors the passage of time. For example, if they can encode four states, $A = 50$ ms, $B = 100$ ms, $C = 200$ ms, and $D = 300$ ms, then the POT is represented by the sequential evolution of the recorder units $A \rightarrow B \rightarrow C \rightarrow D$. Units in the reserve encode 0 ms.

At US onset, this evolution stops and the batch of recorder units is stored in an “archive”. Again, this archive is simply a way of imagining the units after their evolution has been terminated by US onset (unless US onset is too early, for reasons discussed below). We assume that the evolution of the batch is noisy, so that, when the evolution stops, there are some units in each state. Thus, the stored batch is a histogram which constitutes an estimate on the ISI. For example, if the ISI is 200 ms, then there will be some units in states A , B , and C . We choose to imagine the recording medium as consisting of many individual “units”, since we want to use a stored histogram (which displays the frequencies of many elements or samples) to estimate the ISI. The routine of write/OFF \rightarrow write/ON \rightarrow store batch is repeated with each CS presentation, as long as they are spaced far enough apart to respect the refractory timer on the write module. With each presentation, a batch of recorder units is released, evolves and is stored in the archive, updating the total number of recorder units in the archive encoding each time value.

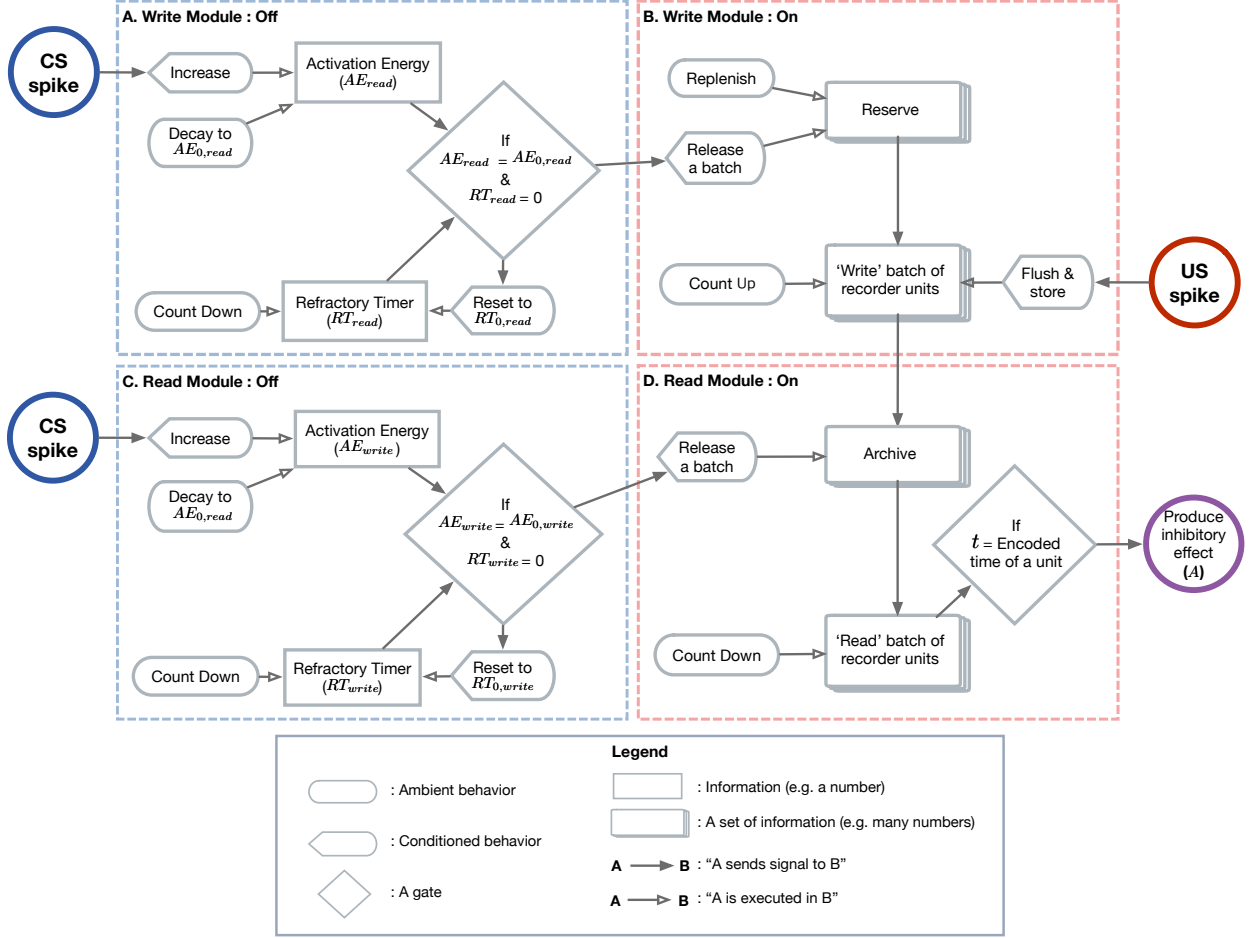


Figure 2: *Model diagram.* This diagram depicts the flow of events leading to the CR in the model Purkinje cell. *Write/OFF*) The write module's switch is controlled by the quantity AE_{write} (upper box). AE_{write} is increased with each CS spike (blue circle) and is otherwise subject to a passive decay. Only when AE_{write} reaches a threshold (diamond) does the cell transition to the write/ON state. When the switch is activated, a refractory timer (lower box) is reset and begins to tick downward until it reaches 0. *Write/ON*) When the write module switches ON, the reserve (tiled boxes, top) continually ejects a batch of recorder units (tiled boxes, bottom). The reserve is then replenished until at maximum capacity. The ejected recorder units advance in state with the passage of time until the first US spike (red circle), at which point the batch is stored in the archive. *Read/OFF*) The switch on the read module behaves exactly as in the write case, but with different parameter values. *Read/ON*) A constant fraction of the archive (tiled boxes, top) is sampled. Each unit "counts down" until the time encoded by its state. At this time (diamond), the unit introduces a fixed inhibitory current to the membrane (purple circle).

Parallel to the writing process is a reading process, which is itself gated by a switch mechanism (Fig. 2, bottom left). Whether the read switch is activated depends on the read module's own activation energy (AE_{read}) and refractory timer. When AE_{read} reaches a threshold, the read module transitions to the ON state (Fig. 2, bottom right). The activated reading module samples a batch of stored recorder units from the archive. Each recorder unit in this batch contributes a fixed amount of inhibition to the cell at a time corresponding to its state. For instance, if batch is released from the archive at time $t = 0$, then a unit in state A will introduce inhibition to the membrane potential at $t = 50$ ms. Note that, the more units there are in a batch encoding a particular time, the greater the inhibition will be at this time. In this way, the shape of the archive's histogram is directly translated into the timecourse of inhibition.

Model details: We model the Purkinje cell as a single-compartment, leaky integrate-and-fire-neuron (Eq. 1) whose parameters are given in Table 1. This simple neural spiking model is appropriate for our purposes, since the Purkinje cell CR is internally generated and therefore does not rely on complicated integration of pre-synaptic spikes:




Figure 3: *ISI and number of trials determine archive shape.* Each curve plots the archive (y -axis, number of units encoding each time step on the x -axis) for different training durations and ISIs. Colors indicate 20 (red), 40 (magenta), 80 (yellow), 160 (green), 320 (blue), 640 (cyan) training trials. Further training trials make the archive grow taller. Panels indicate 200 (top left), 400 (top right), 600 (bottom left) and 800 (bottom right) ms ISIs. Longer ISIs make the archive less peaked. These different shapes translate into different inhibition profiles via the reading module.

$$\tau_m \frac{dV}{dt} = -V(t) + V_{\text{rest}} + R_e CS(t) - R_i A(t) + R_p P(t). \quad (1)$$

The membrane potential is controlled by three sources of current: CS , the 0-1 valued impulse train transmitted by the parallel fibers with resistance R_e ; A , internally-generated inhibitory current the archive with resistance R_e ; and, P , internal, excitatory current generated by the cell's intrinsic pacemaker, which we model as a Poisson impulse train with resistance R_p and rate λ . When the membrane potential reaches a threshold, $V_{\text{threshold}}$, the potential spikes to V_{spike} and then reset to $V_{\text{hyperpolarization}}$.

The switches on both the write and read modules are controlled by a simple exponential decay function involving activation energy :

$$\frac{dAE_{\text{module}}}{dt} = \frac{1}{\tau_{\text{module}}} [AE_{0,\text{module}} - AE_{\text{module}}(t)] + CS(t). \quad (2)$$

Here, AE_{module} refers to the activation energy on the switch for either the write or read module. The time constant for each module's switch, τ_{module} , controls the rate at which the energy passively decays to its resting value, $AE_{0,\text{module}}$, between spikes. The binary impulse train $CS(t)$ increments AE_{module} by 1 with each parallel fiber spike.

If the pre-synaptic spikes arriving at the Purkinje cell from CS stimulation are sufficiently frequent, AE_{write} is driven to a threshold, $AE_{\text{thresh,write}}$ and the write module switches to the ON state. At this point, recorder units are released and allowed to evolve. The write module's reserve can hold a maximum of $R_{\text{max}} = 1$ recorder units, representing 100 % fullness. When the write module switches on, the reserve

Table 1

<i>Parameter</i>	<i>Value</i>
τ_m	5 ms
V_{rest}	- 70 mV
$V_{\text{threshold}}$	-54 mV
$V_{\text{hyperpolarization}}$	-85 mV
V_{spike}	10 mV
R_e	50 Ω
R_i	$2.25 \times 10^6 \Omega$
R_p	50 Ω
τ_{write}	70 ms
τ_{read}	200 ms
$AE_{0,\text{write}}$	0
$AE_{0,\text{read}}$	0
$AE_{\text{thresh},\text{write}}$	2
$AE_{\text{thresh},\text{read}}$	2
R_{max}	1
τ_{reserve}	100 ms
q	1.25×10^{-7} recorder units/ms
σ	40 ms
λ	.3
c	3×10^{-2}

depletes exponentially with time constant τ_{reserve} and its refractory timer is reset so that it cannot detect the CS for a fixed interval. As soon as the reserve begins to empty, it begins to slowly repopulate at a constant rate q . For simplicity, we assume that the state s of a released recorder unit r is real-valued, so that a recorder unit can encode any time between CS onset and US onset up to some maximum granularity Δt . At each time-step, the batch of released recorder units is blurred with a Gaussian kernel of standard deviation $\sigma = 40$ ms to simulate noise.

At the first US spike, the released batch of recorder units is added to a cumulative archive, unless US onset is less than 100 ms. If the first US spike occurs before this time, we assume that the batch is simply discarded since neither the Purkinje cell pause nor the overt blink can be learned for ISIs less than 100 ms. The shape of the batch, and therefore the shape of the archive, is controlled by the number of trials and the ISI (Fig. 3). For short ISIs greater than 100 ms, the stored batch will consist mostly of recorder units approximately encoding the ISI, since noise had little time to take effect. For longer ISIs, noise overwhelms the evolution of recorder units, effectively smoothing the histogram and reducing the peak near the ISI.

The read module acts in parallel to the write module. When pre-synaptic spikes are sufficiently rapid to drive AE_{read} to a threshold, $AE_{\text{thresh},\text{read}}$, the read module switches to the ON state, and a constant fraction c of recorder units are sampled from the archive. After activation, the read module's refractory timer resets so that reading is deactivated for a fixed interval. Each unit in the read batch then introduces a fixed amount of inhibitory current to the membrane at the time corresponding to its state: if read switches ON at time $t = 0$ and the archive contains n_r recorder units encoding state $s = t_1$, then $A(t_1) = cn_r$. Thereby, the shape of the stored archive (i.e the histogram estimate on the ISI) is directly translated into inhibitory current. The recorder units sampled from the archive during reading are permanently removed so that reading will gradually deplete the archive without compensatory addition from writing.

3 Results

Simulation 1: *The cell learns a well-timed pause.* First, we trained the model cell on a 200 ms ISI (Fig. 4a). CS stimulation consisted of a 100 Hz impulse train enduring for 220 ms and US stimulation consisted

of a 500 Hz impulse train enduring for 20 ms and beginning at 200 ms. In this experiment, CS and US stimulation co-terminated. The cell was trained over the course of 400 trials with an ITI of 15 s. We declared that trial on which the average spiking rate during the ISI dropped to below 25 % of the pre-training value to be the start of the CR. In this simulation, the cell learned the pause at trial 177, which took 44.25 minutes of experimental time. Additionally, starting after 300 trials, we removed US stimulation on every 20th trial. On these probe trials (Fig. 4a, red spikes), the cell still produced a pause visually indistinguishable from that of non-probe trials. A peristimulus time histogram, created by binning spikes in windows of 20 ms, averaging across trials, and smoothing with a moving average filter of span 5, is shown in Fig. 4b. The inhibition created by the archive and its effect on the cell's membrane potential are shown in Fig. 4c and 4d respectively.

Next, we trained the cell on a full battery of ISIs ranging from 150 ms to 500 ms in increments of 50

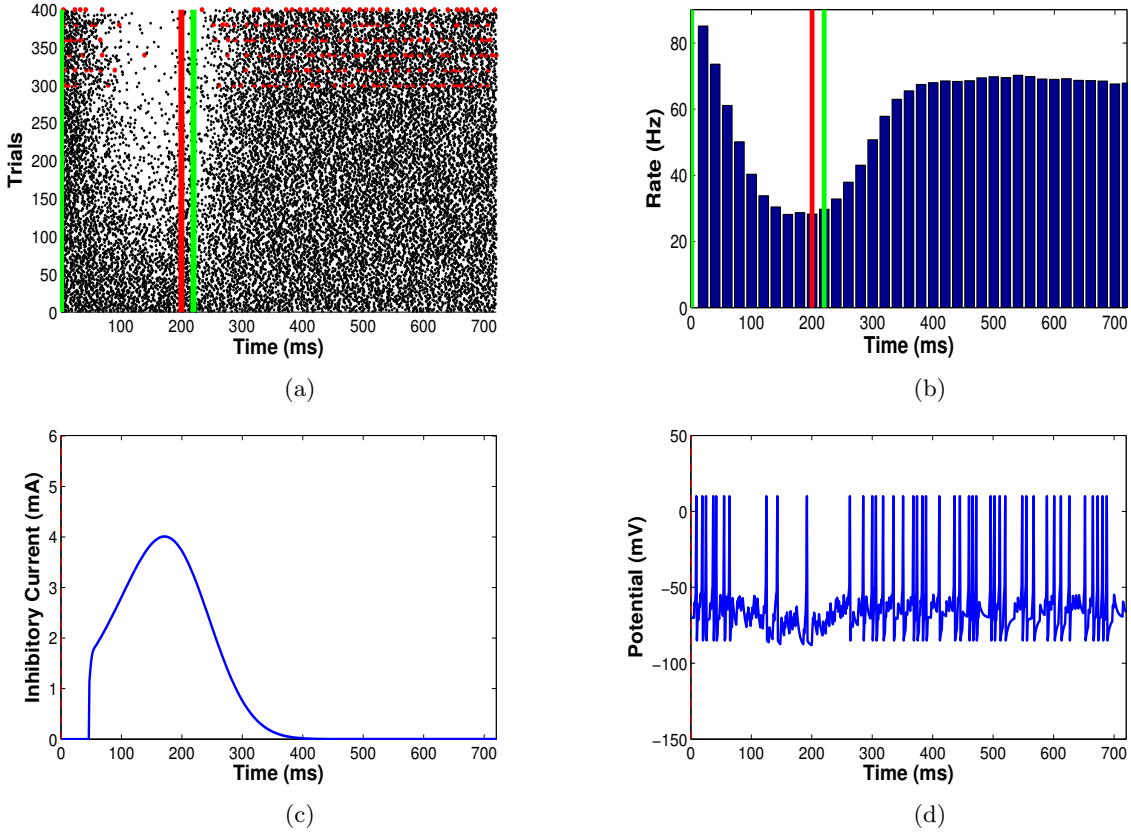


Figure 4: *Basic CR acquisition.* *a)* Raster plot of spike times from Simulation 1. Green lines indicate CS onset and offset. Red line indicates US onset. Probe trials are marked by red spikes. Probe trials still produce the learned pause. *b)* Peristimulus time histogram from Simulation 1. Vertical lines indicate stimulation times, as in panel *a*. *c)* Inhibition caused by the archive on trial 400. Note that the timecourse of inhibition mimics the shape of the archives displayed in Fig. 3a, offset by a few tens of milliseconds. This offset is caused by the time it takes for the read module to detect CS onset and initiate hyperpolarization. *d)* Actual membrane potential during interstimulus period on trial 400. Note the thinning of spikes during the ISI, indicating a well-timed CR

ms. All other training parameters were fixed from the first 200 ms ISI simulation. For this wide range of ISIs, the cell learned pauses with well-timed onsets (Fig. 5a), maxima (Fig. 5b) and offsets (Fig. 5c). On average, pause onsets preceded US onset by 127.5 ms, pause maxima followed US onset by 7.5 ms, and pause offsets followed US onset by 110 ms. The timing of these pause features mimics those found in experimental data (Johansson et al., 2014). By design, the model cannot learn ISIs less than 100 ms.

Simulation 2: *The cell is invariant to probe CS.* Next, we stimulated the trained cell with probe CSs enduring for lengths different from those used during training. Like in Johansson et al. (2014), we trained




Figure 5: *CR temporal features.* In these plots, blue dots indicate a temporal feature of the CR (maxima, onsets, offsets) and the red line indicates US onset *a)* Pause maxima in Simulation 1 occurred essentially at US onset. *b)* Pause onsets always precede the expected US since the archive contains some recorder units encoding low time steps. *c)* Correspondingly, pause offsets always occur shortly after US onset because there are some recorder units encoding time-steps longer than the ISI.

a cell with a 100 Hz 200 ms and probed it with CSs with durations 50 ms and 700 ms (Fig. 6a, 6b) and with frequencies 50 Hz and 200 Hz (Fig. 6c, 6d). Independent of probe CS duration and frequency, the cell still produced a well-timed CR. This invariance arises from the model’s switch mechanism: a small amount of CS stimulation is sufficient to trigger the reading of the stored pause information, during which afferent parallel fiber activity has little effect. This simulation demonstrates that the model satisfies the computational constraint stipulated in the introduction: the learned pause is expressed on probe trials, independent of CS duration.

Simulation 3: *The pause is extinguished over many CS-only trials.* Because the read module gradually depletes the archive, many CS-only trials will eventually extinguish the learned pause response. To demonstrate this behavior, we provided a cell previously trained on a 200 ms ISI with an additional 400 CS-only trials. The raster plot in Fig. 7 shows the gradual extinction of the pause. Within 75 trials, the cell increased to 50 % of its pre-training rate, and by trial 397, it had increased to over 90%. This is in line with experimental data which shows the timescale of extinction is typically less than or equal to that of acquisition.

Simulation 4: *The cell can learn multiple pauses.* Jirenhed et al. (2007) and Johansson et al. (2014) found that interleaving trials with different ISIs teaches the Purkinje cell two CRs so that, given one probe CS, the cell will pause twice. Our model Purkinje cell reproduces this behavior (Fig. 8) in the case of interleaved 200 ms and 500 ms ISI trials with a 30 s ITI. However, in this simulation, both pauses were weaker and took more time to express (197 trials vs. Simulation 1’s 177 trials) using the same detection criterion. This is because only half of the trials contribute to the learning of each ISI, and those trials which do not contribute to one ISI still consume the archive by engaging the read module. As a result, experimental time to acquisition is more than twice that of the original case (98.5 minutes vs. 44.25 minutes), adjusted for ITI.

Simulations 1-4 have been performed in the real Purkinje cell. For simulations 5 and 6, we ran simulations which, to our knowledge, have not been performed experimentally.

Simulation 5: *2 CSs or 2 USs.* First, we trained the model cell with two CSs on each trial. One CS began at $t = 0$ and ended at $t = 100$ ms; another began at $t = 300$ ms and ended at $t = 400$ ms. US onset was at 500 ms and lasted for 20 ms. Hence, the cell was effectively exposed to two ISIs on each trial, 500 ms and 200 ms. Like previous simulations, both CSs were 100 Hz and the US was 500 Hz. The ITI was 15 s. Whereas interleaving two ISIs across trials was found to produce two pauses in Simulation 4, providing the cell with two CSs on each trial produced only one in an experiment lasting 800 trials. (Fig. 9a). The cell only learns the longer ISI since the refractory period on the write switch prevents it from detecting the

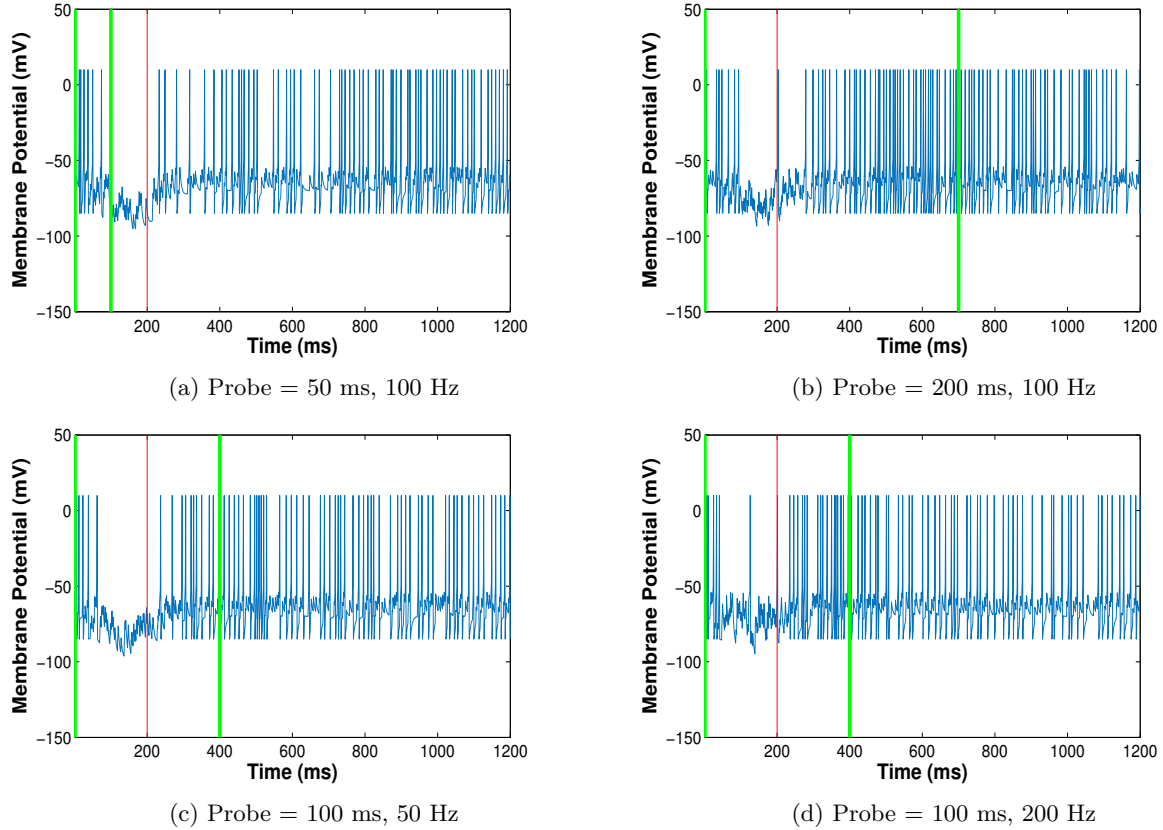


Figure 6: *The CR is invariant to probe duration and frequency.* For all panels, green lines indicate CS onset and offset and red lines indicate US onset. *a, b*) Changing the length of probe CS to 50 or 200 ms does not affect the temporal features of the CR. *c, d*) Changing the frequency of probe CS to 50 or 200 ms does not affect the temporal features of the CR. Note that this simulation also demonstrates the CR is timed to the ISI and not the training CS's duration.

second CS.

Next, we stimulated the cell with two USs on each trial, one at 200 ms and another at 400 ms, both lasting for 20 ms at 500 Hz. The simulation lasted for 400 trials with an ITI of 15 s. The cell learned one pause, timed to the shorter ISI (Fig. 9b). As in the double CS experiment, the cell becomes insensitive to further US stimulation after the first US onset.

Simulation 6: *Trials to acquisition depends on ISI and ITI.* The number of trials until pause acquisition is affected by several factors. First, longer ISIs are harder to learn since noise during writing tends to flatten the histograms added to the archive, which consequently takes more trials to build up enough inhibition to cancel spiking. Short ISIs yield highly peaked histograms and therefore speed up learning. Further, long ITIs tend to speed up learning since they allow the reserve to replenish fully and release a maximum number of recorder units during subsequent reading. Short ITIs, on the other hand, do not allow for full replenishment, so that fewer recorder units are released and the archive grows more slowly.

To investigate how these factors precisely affect our model cell, we ran two parallel experiments. In one experiment, we fixed the ITI at 15 s and trained the cell on ISIs ranging from 100 ms to 1000 ms in steps of 10 ms. For the other experiment, we used the same range of ISIs, but also varied the ITI so that $ITI/ISI = 80$ for each ISI. In the first experiment, trials to acquisition increased monotonically and approximately linearly with ISI (Fig. 10, red line). For short ISIs noise in recorder unit evolution smooths the archive and delays learning. Further, the relatively short ITI prevents the write module's reserve from replenishing fully between trials. However, when the ITI/ISI ratio was held constant, the hindering effect of noise due to increasing ISI was canceled out by the accelerating effect of the increased ITI, resulting in a flat trials-to-




Figure 7: *CR extinction*. Green lines indicate CS onset and offset, and the red dashed line indicates expected US onset, though no US stimulation was provided in this simulation. After many CS-only trials, the CR disappears, though the archive is not empty. It has simply been reduced enough to permit tonic spiking.




Figure 8: *A bimodal pause*. *a)* Green lines indicate CS onset and offset, the red line indicates US onset, and red spikes indicate probe trails. In this simulation, the cell learned two ISIs, one timed to 200 ms and another timed to 500 ms. The CR consists of two pauses with an intervening resumption of spiking between 200 and 400 ms. *b)* The bimodal CR is evident from a raster plot created in the same manner as described in Simulation 1. Presumably, this paradigm could be generalized to more than 2 pauses, as long as the modes of the archive are well-separated enough to allow for spiking resumption between the constituent pauses.

acquisition curve (Fig. 10, blue line). This effect has not yet been tested in the Purkinje cell, though it has been observed in other behavioral experiments (Gibbon, 1977). The finding illustrated in Fig. 10 suggests that trials to acquisition T is given roughly by $T = 8000(\text{ISI}/\text{ITI})$.




Figure 9: *With 2 CSs or USs per trial, only one CR develops.* a) Green lines indicate the onset and offset of the first CS; blue lines indicate the onset and offset of the second CS; red line indicates US onset. A CR appears beginning at around 400 ms and terminates shortly after the US. If the cell had learned the second ISI, the pause would appear and disappear much earlier in the trial, since the second ISI is only 200 ms. b) Red line indicates the first US onset; magenta line indicates second US offset. Spiking resumes around 300 ms since only the shorter ISI is learned.




Figure 10: *ITI affects trials to acquisition.* Red line indicates trials to acquisition in experiment in which ISI varied between 100 and 1000 ms and ITI was fixed at 15 s; blue line indicates trials to acquisition in experiment in which the ISI varied over the same range, but ITI was fixed at 80 times the ISI. The lines intersect when the ITI is 15 s. Otherwise, the red line increases linearly due to the effect of noise on long ISIs and the inability of the reserve to completely replenish in 15 s. The blue line, on the other hand, is constant, since the increased noise of long ISIs is compensated for an increased ITI.

4 Discussion

This purely formal model recapitulates the behavior of the Purkinje cell during eyeblink conditioning. Most importantly, it learns a pause response with critically-timed onset, maximum and offset. Further, the model obeys the two desiderata stipulated in the introduction: the POT representation makes no use of granule cells and is internal to the POT; and, pause learning and expression machinery is gated by a switch. Obeying these criteria saves this model from some of the crucial limitations afflicting previous POT models of the

Purkinje cell. Additionally, the model makes several predictions:

- If the Purkinje cell is stimulated with multiple CSs on each trial, it will only learn the longest ISI.
- If the Purkinje cell is stimulated with multiple USs on each trial, it will only learn the shortest ISI.
- The number of trials until pause acquisition is a constant times the ISI/ITI ratio

Johansson et al. (2014), Chen et al. (2014) and Ryan et al. (2015) may call for a fundamental re-imagining of the neural basis of learning and memory. The model outlined above is a computational proof-of-principle for a learning mechanism inspired by this work, in the case of eye-blink conditioning. Formally, the model in many ways resembles earlier population models of the Purkinje cell learning mechanism, though here a population of neurons has been replaced with a population of abstract “recorder units”. The model makes several straightforward predictions, which can be tested electrophysiologically. Most importantly, a complete understanding of the Purkinje cell will require insight from molecular biology, in order to illuminate the structure of the learning mechanism this paper attempts to formally describe.

References

- Buonomano, D.V. and Mauk, M.D. (1994). Neural Network Model of the Cerebellum: Temporal Discrimination and the Timing of Motor Responses. *Neural Comput.*, 6(1):38–55.
- Cerminara, N.L. and Rawson, J.A. (2004). Evidence that climbing fibers control an intrinsic spike generator in cerebellar Purkinje cells. *J. Neurosci.*, 24(19):4510–7.
- Chen, S., Cai, D., Pearce, K., Sun, P. Y.-W., Roberts, A.C., and Glanzman, D.L. (2014). Reinstatement of long-term memory following erasure of its behavioral and synaptic expression in Aplysia. *Elife*, 3:e03896.
- Eccles, J., Ito, M., and Szentagothai, J. (1967). *The cerebellum as a neuronal machine*. Springer-Verlag, Berlin, Heidelberg, New York.
- Gibbon, J. (1977). Scalar expectancy theory and Weber’s law in animal timing. *Psychol. Rev.*, 84(3):279–325.
- Gormezano, I. and Moore, J. (1969). Classical Conditioning. In Marx, M., editor, *Learn. Process*. Macmillan, New York.
- Halverson, H.E., Khilkevich, A., and Mauk, M.D. (2015). Relating cerebellar purkinje cell activity to the timing and amplitude of conditioned eyelid responses. *J. Neurosci.*, 35(20):7813–32.
- Harvey, R.J. and Napper, R.M. (1991). Quantitative studies on the mammalian cerebellum. *Prog. Neurobiol.*, 36(6):437–63.
- Heiney, S.A., Kim, J., Augustine, G.J., and Medina, J.F. (2014). Precise control of movement kinematics by optogenetic inhibition of Purkinje cell activity. *J. Neurosci.*, 34(6):2321–30.
- Hesslow, G. (1994a). Correspondence between climbing fibre input and motor output in eyeblink-related areas in cat cerebellar cortex. *J. Physiol.*, 476(2):229–44.
- Hesslow, G. (1994b). Inhibition of classically conditioned eyeblink responses by stimulation of the cerebellar cortex in the decerebrate cat. *J. Physiol.*, 476(2):245–56.
- Hesslow, G. (1995). Classical conditioning of eyeblink in decerebrate cats and ferrets. In Ferrell, W. and Proske, U., editors, *Neural Control Mov.*, pages 117–22. Plenum Press, New York.
- Hesslow, G. and Ivarsson, M. (1994). Suppression of cerebellar Purkinje cells during conditioned responses in ferrets. *Neuroreport*, 5(5):649–52.
- Hesslow, G., Svensson, P., and Ivarsson, M. (1999). Learned movements elicited by direct stimulation of cerebellar mossy fiber afferents. *Neuron*, 24(1):179–85.

- Hesslow, G. and Yeo, C. (2002). The Functional Anatomy of Skeletal Conditioning. In Moore, J., editor, *A Neurosci. Guid. to Class. Cond.*, pages 86–146. Springer-Verlag, New York.
- Jirenhed, D.-A., Bengtsson, F., and Hesslow, G. (2007). Acquisition, extinction, and reacquisition of a cerebellar cortical memory trace. *J. Neurosci.*, 27(10):2493–502.
- Jirenhed, D.-A. and Hesslow, G. (2011a). Learning stimulus intervals—adaptive timing of conditioned purkinje cell responses. *Cerebellum*, 10(3):523–35.
- Jirenhed, D.-A. and Hesslow, G. (2011b). Time course of classically conditioned Purkinje cell response is determined by initial part of conditioned stimulus. *J. Neurosci.*, 31(25):9070–4.
- Johansson, F., Carlsson, H., Rasmussen, A., Yeo, C.H., and Hesslow, G. (2015). Activation of a temporal memory in Purkinje cells by the mGluR7 receptor. *Press*.
- Johansson, F. and Hesslow, G. (2014). Theoretical considerations for understanding a Purkinje cell timing mechanism. *Commun. Integr. Biol.*, 7(6):e994376.
- Johansson, F., Jirenhed, D.-A., Rasmussen, A., Zucca, R., and Hesslow, G. (2014). Memory trace and timing mechanism localized to cerebellar Purkinje cells. *Proc. Natl. Acad. Sci. U. S. A.*, 111(41):14930–4.
- Kehoe, E. and Macrae, M. (2002). Fundamental Behavioral Methods and Findings in Classical Conditioning. In Moore, J., editor, *A Neurosci. Guid. to Class. Cond.*, pages 171–231. New York.
- Mauk, M.D., Steinmetz, J.E., and Thompson, R.F. (1986). Classical conditioning using stimulation of the inferior olive as the unconditioned stimulus. *Proc. Natl. Acad. Sci. U. S. A.*, 83(14):5349–53.
- Medina, J.F., Garcia, K.S., Nores, W.L., Taylor, N.M., and Mauk, M.D. (2000). Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *J. Neurosci.*, 20(14):5516–5525.
- Medina, J.F. and Mauk, M.D. (2000). Computer simulation of cerebellar information processing. *Nat. Neurosci.*, 3 Suppl:1205–11.
- Mostofi, A., Holtzman, T., Grout, A.S., Yeo, C.H., and Edgley, S.A. (2010). Electrophysiological localization of eyeblink-related microzones in rabbit cerebellar cortex. *J. Neurosci.*, 30(26):8920–34.
- Ryan, T.J., Roy, D.S., Pignatelli, M., Arons, A., and Tonegawa, S. (2015). Engram cells retain memory under retrograde amnesia. *Science (80-.)*, 348(6238):1007–1013.
- Steinmetz, J.E., Lavond, D.G., and Thompson, R.F. (1989). Classical conditioning in rabbits using pontine nucleus stimulation as a conditioned stimulus and inferior olive stimulation as an unconditioned stimulus. *Synapse*, 3(3):225–33.
- Svensson, P. and Ivarsson, M. (1999). Short-lasting conditioned stimulus applied to the middle cerebellar peduncle elicits delayed conditioned eye blink responses in the decerebrate ferret. *Eur. J. Neurosci.*, 11(12):4333–40.
- Wetmore, D.Z., Jirenhed, D.-A., Rasmussen, A., Johansson, F., Schnitzer, M.J., and Hesslow, G. (2014). Bidirectional plasticity of Purkinje cells matches temporal features of learning. *J. Neurosci.*, 34(5):1731–7.
- Yamazaki, T. and Tanaka, S. (2009). Computational models of timing mechanisms in the cerebellar granular layer. *Cerebellum*, 8(4):423–32.
- Yeo, C.H., Hardiman, M.J., and Glickstein, M. (1985). Classical conditioning of the nictitating membrane response of the rabbit. I. Lesions of the cerebellar nuclei. *Exp. brain Res.*, 60:87–98.